

METHOD AND APPARATUS FOR PROVIDING INTERNAL TABLE EXTENSIBILITY BASED ON PRODUCT CONFIGURATION

Inventors:

Gaurav Singh

Ali Kani

Kiran Kattel

Sridhar Subramanian

Brian Hang Wai Yang

FIELD

[0001] The invention relates generally to the field of networking, and more particularly to lookup table configuration control.

BACKGROUND

[0002] In networking systems, routers and/or switches typically use lookup tables to facilitate or increase the speed of the forwarding of packets of information. Such tables can be implemented in hardware, for example, as content-addressable memory (CAM), which can include static random access memory (SRAM) type memory cells. Among the types of forwarding that can be accommodated in typical systems are Layer 2 (L2) and/or Layer 3 (L3) forwarding. L2 tables can include Media Access Control (MAC) address lookup and L3 tables can include Internet Protocol (IP) address lookup. Other example table types include IP MultiCast (MC), IP Next Hop Table (NHT), and IP Longest Prefix Match (LPM) tables. Internal (i.e., on-chip) tables can be used for the L2/L3 forwarding and these tables can, in many cases, be modified by hardware and/or software control.

[0003] For systems that support more than one type of lookup, a typical approach for table management is shown in FIG. 1A. FIG. 1A is a block diagram of a conventional L2/L3 lookup table arrangement used for both L2 and L3 type lookups. General reference character 100 indicates Router/Switch 102 includes L2 Controller 104 for control of L2 Table 108 as well as L3 Controller 106 for control of L3 Table 110. L2 Table 108 has an entry width of 72-bits and L3 Table 110 has an entry width of 45-bits. In this conventional approach, one table is always configured as L2 type and the other is always configured as L3 type.

[0004] Referring now to FIG. 1B, a block diagram of a conventional L2/L3 lookup table arrangement used for only L2 type lookups is shown and indicated by the general reference character 150. Like numbers are typically used herein to indicate the same or similar

components as in the diagram of FIG. 1A. In FIG. 1B, L3 Controller 106 has switched control from “enable” to “disable” so that L3 Table 110 is not used for L2 type lookups. Thus, the table designated for L3 lookups goes unused, as indicated by the cross-hatching in FIG. 1B, in this conventional approach. This can represent a large portion of the die area that is not utilized. Accordingly, this approach is not cost effective.

[0005] Consequently, what is needed is a solution that can control and allocate internal lookup tables so that each table is substantially utilized for a variety of different product configurations.

SUMMARY

[0006] The invention overcomes the identified limitations and provides a solution that can control and allocate internal lookup tables so that each table is substantially utilized for a variety of different product configurations.

[0007] According to embodiments, a configurable lookup table system includes a first controller coupled to a first lookup table and a second controller coupled to a second lookup table. The first controller configures the first lookup table for a first type lookup, which can be a Layer 2 or Media Access Control (MAC) type. The second controller configures the second lookup table based on a mode determination. If in a first mode, the second lookup table can be configured for a second type lookup, which can be a Layer 3 or Internet Protocol (IP) type. If in a second mode, the second lookup table can be configured for the first type lookup.

[0008] According to another aspect of the embodiments, a method of controlling a lookup table system includes determining a first mode or a second mode. If a first mode, the method includes using a first lookup table for a first type lookup and using a second lookup table for a second type lookup. If a second mode, the method includes splitting the first type lookup into a first level and a second level, using the second lookup table for the first level, and using the first lookup table for the second level.

[0009] This approach provides an efficient scheme for controlling and using multiple internal lookup tables for a variety of product configurations.

BRIEF DESCRIPTION OF THE FIGURES

[0010] Embodiments of the invention are described with reference to the FIGS, in which:

[0011] FIG. 1A is a block diagram of a conventional L2/L3 lookup table arrangement used for both L2 and L3 type lookups.

[0012] FIG. 1B is a block diagram of a conventional L2/L3 lookup table arrangement used for only L2 type lookups.

[0013] FIG. 2A is a block diagram of an L2/L3 lookup table arrangement in a first mode according to an embodiment.

[0014] FIG. 2B is a block diagram of an L2/L3 lookup table arrangement in a second mode according to an embodiment.

[0015] FIG. 3 is a flow diagram of a method of processing lookups according to an embodiment.

DETAILED DESCRIPTION

[0016] Embodiments of the invention are described with reference to specific diagrams depicting system arrangements and methods. Those skilled in the art will recognize that the description is for illustration and to provide the best mode of practicing the invention. The description is not meant to be limiting. For example, reference is made to several types of forwarding lookup tables, including IP and MAC types, but the invention is applicable to other types of lookup tables as well. Further, table organization, including separate tables and the like, merely provide example implementations and should not be construed as limiting.

[0017] A router/switch chip according to embodiments can work in different modes. In one of the modes, the chip can perform both Layer 2 (L2) and Layer 3 (L3) type forwarding. This utilizes both the L2 and L3 tables included on the die. In another mode, the chip can re-configure so as to perform only L2 type forwarding. In this mode, L3 type tables are not needed. Rather than not use these tables in such a mode, as is commonly done in conventional approaches, a mechanism described herein can logically increase the size of the L2 type table by utilizing the L3 table. According to embodiments, L3 tables can be designed in a format that encompasses both L2 and L3 type entries. Further, a unified memory controller can be employed for both modes.

[0018] Referring now to FIG. 2A, a block diagram of an L2/L3 lookup table arrangement in a first mode according to an embodiment is shown and indicated by the general reference character 200. Router/Switch 202 can include L2 Controller 204 for control of L2 Table 208 and Unified L2/L3 Controller 206 for control of L3 Table 210. In this one very particular example arrangement in an operational mode accommodating both L2 and L3 type

lookups, a table entry width for L2 Table 208 may be 72-bits and a table entry width for L3 Table 210 may be 45-bits. TABLE 1 and TABLE 2 below show example bit designations for entries according to such example lookup types.

TABLE 1 (L2 Table Entry):

Field:	Width (bits):	Description:
MAC Address	A	MAC destination address of the next hop
VLAN	B	VLAN ID to be used
Port Number	C	Exit port number (may be consistent with L2 table encoding)
Parity	D	Entry parity check bits
Total	72	

TABLE 2 (L3 Table Entry):

Field:	Width (bits):	Description:
Valid	E	Entry valid indication
Compare Key	F	Address key used for compare
Parity	G	Entry parity check bit
Total	45	

[0019] Referring now to FIG. 2B, a block diagram of an L2/L3 lookup table arrangement in a second mode according to an embodiment is shown and indicated by the general reference character 250. Like numbers are typically used herein to indicate the same or similar components as in the diagram of FIG. 2A. According to embodiments, in this second operational mode, Unified L2/L3 Controller 206 can configure the L3 type table to function as an L2 type table. This can be done by configuring the table into two effective portions as shown: a first portion 210A and a second portion 210B. The first portion 210A can include usable L2 type entry allocation and the second portion 210B can be a reserved or unused portion, depending on the application. In this fashion, sets of two 45-bit wide entries can be effectively or logically combined to form an entry width of up to 90-bits, as one example. Accordingly, the general format of an L3 type table portion can be made to accommodate both L2 and L3 type lookups. Further, in this second mode, a hash table approach can be used for a “split” lookup where the 210A portion can form a “Level 1” and L2 Table 208 can form a

“Level 2” indexed by the hash table. In this approach, a reference of perhaps a few bits can be done and then indexed into one or both of the example tables. Because of the possibility of finding a match or “hit” in each table, corresponding to two different addresses mapping to the same MAC address, there should be a checking mechanism and a compare performed to settle on the correct result.

[0020] TABLE 3 below shows example bit designations for entries in an L3 type table reconfigured to support L2 type lookups, according to embodiments.

TABLE 3:

Field:	Width (bits):	Description:
MAC Address	H	MAC destination address of the next hop
VLAN	I	VLAN ID to be used
Port Number	J	Exit port number (may be consistent with L2 table encoding)
Parity	K	Entry parity check bits
Reserved	L	
Total	90	

[0021] Referring now to FIG. 3, a flow diagram of a method of processing lookups according to an embodiment is shown and indicated by the general reference character 300. First Mode decision step 302 can determine whether the product is configured for both L2 and L3 type lookups (first mode) or only L2 type lookups (second mode), for example. If in the first mode, the lookup tables used depends on the designated lookup type, as determined in L2 Type Forwarding 306 decision box. If it is an L2 type lookup, the flow proceeds to Use First Lookup Table 310 and then to Complete Lookup 312. If it is an L3 type lookup, the flow proceeds to Use Second Lookup Table 308 and then to Complete Lookup 312. If in the second mode, the flow enters Split Lookup 304 function block. Two paths can exit this box: a “Level 1” path to Use Second Lookup Table 308 and a “Level 2” path to Use First Lookup Table 310. Alternatively, the “Level 1” and “Level 2” search functions can be done in a serial fashion, as opposed to the parallel functionality shown in FIG. 3. For example, a Level 1 lookup may be performed and, if a miss indication, a Level 2 lookup may then be performed. If a hit indication was obtained in the Level 1 lookup, then the serial flow may not proceed to a Level 2 lookup. Whether a serial or a parallel Level 1 and Level 2 arrangement is used may depend

on system design considerations. In either arrangement, using the methods according to embodiments, a product configured for both L2 and L3 lookup types or only for L2 lookup types can efficiently utilize available internal lookup tables.

[0022] Advantages of embodiments of the invention include the control of lookup tables to substantially utilize the implemented internal tables based on a variety of possible product configurations.

[0023] Having disclosed exemplary embodiments and the best mode, modifications and variations may be made to the disclosed embodiments while remaining within the subject and spirit of the invention as defined by the following claims.